

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND APPARATUS FOR DYNAMIC BANDWIDTH MANAGEMENT
FOR VOICE TRAFFIC IN A DIGITAL COMMUNICATIONS NETWORK**

INVENTOR:

BRUCE BUFFAM

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 WILSHIRE BOULEVARD, SEVENTH FLOOR
LOS ANGELES, CA 90025-1026

(408) 720-8300

"Express Mail" mailing label number: EL627469805US

Date of Deposit: March 30, 2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Geneva Walls

(Typed or printed name of person mailing paper or fee)

Geneva Walls
(Signature of person mailing paper or fee)

March 30, 2001

(Date signed)

METHOD AND APPARATUS FOR DYNAMIC BANDWIDTH MANAGEMENT FOR VOICE TRAFFIC IN A DIGITAL COMMUNICATIONS NETWORK

FIELD OF THE INVENTION

[0001] The present invention relates to managing data sent through a digital network and, more specifically, to a method and apparatus for dynamic bandwidth management.

BACKGROUND OF THE INVENTION

[0002] Asynchronous Transfer Mode (ATM) or “cell switching” is a method of transmitting digital information wherein the information is broken into equal size units named “cells”. Each cell is 53 octets in length that consists of 5 octets of cell header information and 48 octets of payload. Using the portion of the cell header known as the virtual path index (VPI) and the Virtual Connection Index (VCI) multiple streams of user data can be carried over a single ATM physical media. Most ATM switches support a wide range of media types for digital transmission such as T1, E1, DS-3, E-3, OC-3, OC-12, etc. Regardless of media type, each media interface within a given ATM switch is identified uniquely and thus each user stream entering or leaving the switch can be identified uniquely using the following naming convention, <interface_index.virtual_path_index.virtual_channel_index>. Each unique user stream within a given media interface is known as a connection segment. A complete ATM connection is the concatenation of one or more ATM connection segments. The act of concatenating the connection segments is the primary function of an ATM switch. There are two forms of concatenation, a unicast cross-connection and a

multicast cross-connection. A unicast cross connection is described using the following naming pair, <>upstream:interface_index.virtual_path_index.

virtual_channel_index>,<downstream:interface_index.virtual_path_index.virtual_channel_index>. A multicast cross-connect differs slightly from the above

5 example in that there can be more than one downstream connection segment.

[0003] Voice transport over any network technology poses a unique set of problems or requirements for the underlying networking technology. Voice is extremely intolerant of delay variation. Excessive “jitter” in the form of queuing delays or payload disordering render it difficult to accurately recreate the voice

10 conversation at the receiving end. Excessive total delay caused by cumulative switching delays and large sampling intervals can lead to excessive total delay resulting in excessive echo. Loss of voice payloads is tolerated providing the loss detection techniques do not introduce increased “jitter”. Similar to ATM, an ideal solution would divide voice and data into uniform chunks for carriage

15 across the core of the network and introduce queuing mechanisms to ensure constant rate delivery of the voice payloads. One issue with the use of ATM for the transport of voice is the size of the payload. Prior to the advent of ATM AAL2 the predominant techniques for voice transmission, namely AAL1 and AAL5 proved to be either wasteful of network bandwidth because most voice

20 payloads are much smaller than 48 octets or wasteful of bandwidth because of excessive header and trailer overhead. These issues are particularly exacerbated over low speed ATM access such as DSL or T1/E1 transmission medium. For example, a T1 interface running TDM can accommodate 24

channels of uncompressed voice whereas 24 channels of uncompressed voice over ATM will not fit in a T1 cell interface.

[0004] These issues lead to the development of ATM Adaptation Layer 2(AAL2) . In AAL2, another layer of stream identification known as the Channel

- 5 Identifier (CID) was introduced and a format devised to allow the packing of multiple voice streams sharing a common destination into a single ATM connection. Using this technique a voice stream is uniquely identified using the following naming convention,

<interface_index.virtual_path_index.virtual_channel_index.channel_id>. This

- 10 opens the opportunity to provide a switching function for voice flows analogous to ATM cell switching such that a voice connection could consist of one or more concatenated voice cross-connects, described as follows, <<upstream:
interface_index.virtual_path_index.virtual_channel_index.channel_id>,<downstre
am:interface_index.virtual_path_index.virtual_channel_index.channel_id>>.

- 15 AAL2 provides for the setup of ATM AAL2 connections that are configured to support up to 255 voice channels carried within or can be restricted to just one voice channel carried within. AAL2 connections can be configured as permanent virtual circuits by the operator or they can also be set up on demand as switched virtual circuits using the ATM networking layer. Given that AAL2 implements a
20 voice network on top of an ATM network infra-structure using an overlay technique, it is natural to expect that the configuration of multiplexed connections between any two voice switching nodes involves some network engineering analysis by the operator. Inevitably this analysis will raise questions regarding

how much bandwidth to assign to the multiplexed connections, what action should be taken when bandwidth is exhausted on a multiplexed connection and how to promote efficient bandwidth utilization within these connections. An over-commitment of bandwidth can reduce network bandwidth efficiency and an under-commitment can deny service to users, particularly at peak times.

Particularly, in cases where the underlying ATM network has plenty of unused capacity, service denial constitutes an error and perhaps a failure to honor the service level agreement (SLA) between the network operator and the customer.

SUMMARY OF THE INVENTION

- [0005]** A method and apparatus for dynamic bandwidth management for voice traffic in a digital communications network is disclosed. In one embodiment, a method for use in a digital communications network comprises
- 5 checking a multiplexed connection's bandwidth capacity to carry a call over a link; and overflowing the call onto a non-multiplexed connection, when the multiplexing connection's bandwidth is insufficient to carry the call.
- [0006]** Other features and advantages of the present invention will be apparent from the accompanying drawings and from the detailed description that
- 10 follows below.

USPTO Patent Number 81862 P211

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The present invention is illustrated by way of example, and not limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements and in which:

5

[0008] **Figure 1** illustrates an ATM network with dynamic bandwidth management for high priority traffic in a digital communications network;

[0009] **Figure 2** illustrates exemplary Q.AAL2 data cells as implemented

10 over the structure of ATM cell relay;

[0010] **Figure 3** illustrates the voice switching component of frame relay card; and

15 [0011] **Figure 4** illustrates an exemplary flow diagram of the process performed by switching component.

DETAILED DESCRIPTION

[0012] A method and apparatus for dynamic bandwidth management for

voice traffic in a digital communications network is disclosed. As will be

discussed in greater detail below, in one embodiment, a method for use in a

5 digital communications network comprises checking a multiplexed connection's bandwidth capacity to carry a call over a link; and overflowing the call onto a non-multiplexed connection, when the multiplexing connection's bandwidth is insufficient to carry the call.

[0013] Typically, high-speed data switches are capable of handling many

10 different classes of cell traffic, each class having different characteristics and

different service requirements. The various classes of cell traffic might include

high priority traffic, voice, high-speed deterministic traffic, bursty data, etc. For

example, voice traffic is relatively delay sensitive and insensitive to occasional

cell loss. In contrast, data traffic, such as file transfers, is relatively insensitive to

15 delay but is data loss sensitive.

[0014] **Figure 1** illustrates an ATM network with dynamic bandwidth

management for high priority traffic in a digital communications network.

Network 100 includes switch 110, switch 115, and switch 170. Switches 110,

115, and 170 may be high-speed switches such as those manufactured by Cisco

20 Systems of San Jose, CA. Included in switches 110, 115, and 170 are frame

relay cards 180 having a voice-switching component for dynamically managing

bandwidth in an ATM trunk.

[0015] A trunk 130 connects switches 110 and 170. Trunk 130 includes non-multiplexing channel 150 and multiplexing channel 140. Channels 140, 150 are both bi-directional. Each channel 140, 150 may carry thousands of connections or calls. Multiplexed channel 140 may be a multiplexed Q.AAL2 channel that carries time sensitive data, such as voice or video calls. Non-multiplexed channel 150 may be a non-multiplexed Q.AAL2 channel that carries digital data that is not time or frequency sensitive.

5 [0016] Non-multiplexing channel 155 and multiplexing channel 145 connect switches 170 and 115. Channels 145 and 155 have the same properties as channels 140 and 150 as described above. However, non-multiplexing channel 145 has all of its available bandwidth in use.

10 [0017] Cells 160 are multiplexed over channels 140 and 145. Cells 160 may be Q.AAL2 cells. **Figure 2** illustrates exemplary Q.AAL2 data cells 160 as implemented over the structure of ATM cell relay. A Q.AAI2 cell 160 may be composed of AAL2 Convergence Packet Sub Layer (CPS) packets 210. Each CPS packet 210 has a three byte header 220 containing an 8 bit CID, a 5 bit header error correction field (HEC), a 6 bit length field and a 5 bit User to User Indication (UUI) field. CPS packets 210 have a maximum payload 230 of 45/64 octets of data.

15 [0018] The packets 210 that have a common CID are combined into a Protocol Data Unit (PDU) 240. PDU 240 contains a unique header 250 and payload 260. PDU 240 is broken into multiple payloads 241-244. Each payload 241-244 is stored within cells 160. Each cell is associated with a unique

interface, virtual connection index (VCI) and virtual path index (VPI). Although only eight CPS packets 210 are shown, more or less may be handled.

[0019] Referring back to **Figure 1**, multiple cells 160 are shown on channels 140, 145, 150, and 155. CID=x cell 160 on multiplexing channel 140 is a new call flowing into switch 170 from switch 11 as indicated by the arrow 190. Switch 170 attempts to route CID=x cell 160 to switch 130 (as shown by arrow 195), but finds that multiplexing channel 145 has no available bandwidth. Overflow occurs, and switch 170 adds a non-multiplexing connection 155 to carry CID=x cell 160.

10 [0020] **Figure 3** illustrates the voice-switching component 300 of frame relay card 180. One embodiment of voice-switching component 300 comprises a system bus 320 for communicating information, and a processor 310 coupled to bus 320 for receiving messages from line cards 180 and terminating commands carried in the messages. Voice-switching component 300 further 15 comprises a random access memory (RAM) or other dynamic storage device 325 (referred to herein as main memory), coupled to bus 320 for storing information and instructions to be executed by processor 310. Main memory 325 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor 310. Voice-switching 20 component 300 also may include a read only memory (ROM) and/or other static storage device 326 coupled to bus 320 for storing static information and instructions used by processor 310.

[0021] A data storage device 327 such as a magnetic disk or optical disc and its corresponding drive may also be coupled to voice-switching component 300 for storing information and instructions. Switching component 300 includes buffer 340 for storing CPS packets 210 until cell 160 is filled to capacity and 5 transmitted across channel 150. The buffer 340 may also transmit a cell at set time periods. Buffer 340 is also connected to bus 320.

[0022] Although embodiments of the present invention are described as having both software and hardware elements, alternative embodiments may be all hardware, all software, or a combination of each. The software implementing 10 the present invention can be stored in RAM 325, a mass storage device available through disk interface, or other storage medium accessible to CPU 310. This software may also be resident on an article of manufacture comprising a computer usable mass storage medium or propagated digital signal having computer readable program code embodied therein and being readable by the 15 mass storage medium and for causing CPU 310 to control tasks on networking cards in accordance with the teachings herein.

[0023] **Figure 4** illustrates an exemplary flow diagram of the process 400 performed by switching component 300. The process starts at block 401. Flow continues to processing block 410, where switching component 300 monitors 20 the bandwidth usage on multiplexing channel 150. At decision block 420, Switching component 300 determines if there is sufficient bandwidth available on channel 150 to add a new call. If there is sufficient bandwidth available, flow

continues to processing block 430. At processing block 430, the new call is routed over multiplexing channel 150 and the process completes at block 499.

[0024] If there is insufficient bandwidth available on channel 150, then flow continues to processing block 440. At processing block 440, a new single non-

5 multiplexing connection is added to channel 140 for the new call. Flow continues to processing block 450 where the new call is routed over the new non-multiplexing connection established on channel 140. In one embodiment, the new connection is established by presenting calls that overflow the bandwidth of the multiplexing channel 150 to the ATM Q.2931 layer and signaling the call as a
10 non-multiplexed AAL2 connection. Flow continues to processing block 460 where the non-multiplexing connection is torn down once the new call is complete. The process ends at block 499.

[0025] In one embodiment, any cells that overflow from channel 150 to channel 140 may return to a multiplexing channel once bandwidth becomes

15 available on the connection.

[0026] In another embodiment, the present techniques avoid rejecting calls and allow a user to configure a multiplexed connection with sufficient bandwidth to satisfy an average daily demand. Peak or high volume periods may be satisfied using the overflow technique.

20 **[0027]** In the forgoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the

appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive manner.